

Listing of Claims

1-3. (Canceled)

4. (Previously Presented) A method of reducing the number of times a tree data structure is rebalanced comprising the steps of:

- (I) allowing a sub-tree of the tree data structure to grow until a number of unbalanced levels reaches a threshold greater than one; and
- (II) rebalancing the tree data structure when the threshold is reached by:
 - (a) developing first and second sets of rebalancing operation tasks, the first set of operation tasks operable to effect a first set of element state transitions and the second set of operation tasks operable to effect a second set of element state transitions, the first and second set of element state transition being distinct one from the other;
 - (b) performing the first set of operation tasks in a first phase; and
 - (c) performing the second set of operation tasks in a second phase.

5-9. (Canceled)

10. (Original) A method of performing a rebalancing operation upon a tree data structure comprising the steps of:

- (a) allowing a sub-tree of the tree data structure to grow unbalanced to a threshold level greater than one;

- (b) developing, in the case where the sub-tree reaches the threshold level, first and second sets of rebalancing operation tasks, the first and second set of rebalancing operation tasks operable to effect a first and second set of element state transitions respectively;
- (c) performing the first set of operation tasks in a first phase; and
- (d) performing the second set of operation tasks in a second phase.

11-14. (Canceled)

15. (Currently Amended) A ~~system~~ computer readable storage medium having embodied thereon computing instructions executable by a processor to perform a method, the method comprising:

~~means for~~ storing a tree data structure;

~~means for~~ tracking the execution of operations upon the tree data structure; and

~~means for~~ rebalancing the tree data structure when an unbalanced sub-tree of the tree data structure reaches a threshold level greater than one, the rebalancing including a first rebalancing phase in which rebalancing operations are executed in parallel and nodes of the unbalanced sub-tree are unlocked, and a second rebalancing phase in which different rebalancing operations are executed.

16. (Canceled)

17. (Previously Presented) A method of deferring the rebalancing of a tree data structure comprising the steps of:

- (a) tracking the performance of operations upon the tree data structure; and
- (b) rebalancing the tree data structure when an unbalanced sub-tree of the tree data structure reaches a threshold level greater than one, the rebalancing further comprising creating a first set of rebalancing operation tasks, the first set of rebalancing operation tasks being characterized by navigation of the tree data structure using at least an existing link, creating a second set of rebalancing operation tasks, the second set of rebalancing operation tasks being different from the first set of rebalancing operation tasks and being characterized by location of elements within the tree data structure using at least one pointer external to the tree data structure and created by the first set of rebalancing operation tasks, and performing at least one operation task of the first set of rebalancing operation tasks in a first phase and at least one of the second set of rebalancing operation tasks in a second phase.

18. (Original) A method of deferring the rebalancing of a tree data structure comprising the steps of:

- (a) tracking the performance of operations upon the tree data structure; and
- (b) rebalancing the tree data structure when an unbalanced sub-tree of the tree data structure reaches a threshold level greater than one, the rebalancing further comprising executing simultaneous rebalancing operations on the tree data

structure including performing any first phase operation task of each of the simultaneous rebalancing operations in a first phase using parallel processes, developing a set of serial rebalancing operations during the first phase, and performing any second phase operation task of each of the simultaneous rebalancing operations in a second phase, the second phase operation task having at least one of the set of serial rebalancing operations.

19. (Previously Presented) A method of rebalancing a tree data structure, the method comprising:
- allowing a sub-tree of the tree data structure to grow unbalanced until a threshold level is reached;
 - developing a first set of rebalancing operation tasks, the first set of operation tasks operable in parallel on one or more unlocked nodes of the tree data structure during a first phase of the rebalancing;
 - developing a second set of rebalancing operation tasks during execution of the first set of rebalancing operation tasks; and
 - executing the second set of rebalancing operation tasks during a second phase of the rebalancing.
20. (Previously Presented) The method of claim 19, wherein execution of the second set of rebalancing operation tasks is performed without navigating between nodes of the sub-tree.

21. (Previously Presented) The method of claim 19, wherein execution of the first set of rebalancing operations includes generating a list of pointers to nodes requiring updating in the second phase of rebalancing.
22. (Previously Presented) A method of rebalancing a tree data structure, the method comprising:
allowing a sub-tree of the tree data structure to grow unbalanced until a threshold level is reached;
executing a first set of rebalancing operation tasks during a first rebalancing phase, the first rebalancing phase being characterized by navigation between nodes of the sub-tree; and
executing a second set of rebalancing operation tasks during a second rebalancing phase, the second rebalancing phase including navigation to two or more nodes of the sub-tree, the navigation being independent of pointers between nodes of the sub-tree.
23. (Previously Presented) The method of claim 22, wherein the first set of rebalancing operation tasks are performed on unlocked nodes of the sub-tree.
24. (Previously Presented) The method of claim 22, wherein the first set of rebalancing operation tasks includes a plurality of operation tasks configured for parallel execution.
25. (Previously Presented) A method of maintaining a tree data structure, the method comprising:

allowing the tree data structure to grow unbalanced;

performing a first set of rebalancing operation tasks during a first rebalancing phase on a plurality of nodes in the tree data structure, the first set of rebalancing operation tasks being configured for execution while the plurality of nodes are unlocked and for insertion and deletion of nodes; and

performing a second set of rebalancing operation tasks on the plurality of nodes in a second rebalancing phase, the second set of rebalancing operation tasks being different than the first set of rebalancing operation tasks and being configured for further operations on the plurality of nodes, the second rebalancing phase occurring after completion of the first rebalancing phase.

26. (Previously Presented) The method of claim 25, wherein the first set of rebalancing operation tasks are performed in parallel.

27. (Previously Presented) A method comprising:

storing a tree data structure;

tracking the execution of operations upon the tree data structure; and

rebalancing the tree data structure when an unbalanced sub-tree of the tree data structure reaches a threshold level greater than one, the rebalancing including a first rebalancing phase in which rebalancing operations are executed in parallel and nodes of the unbalanced sub-tree are unlocked, and a second rebalancing phase in which different rebalancing operations are executed.

28. (Previously Presented) The method of claim 10, wherein the first and second set of element state transitions each include changing pointers to nodes of the tree data structure.

29. (New) A system for deferring the rebalancing of a tree data structure comprising:

- (a) a memory for storing the tree data structure;
- (b) a processor coupled to the memory, the processor operable to track the performance of operations upon the tree data structure and rebalance the tree data structure when an unbalanced sub-tree of the tree data structure reaches a threshold level greater than one, the rebalancing including a first rebalancing phase in which rebalancing operations are executed in parallel and nodes of the unbalanced sub-tree are unlocked, and a second rebalancing phase in which different rebalancing operations are executed.